

Variational Boosting: Iteratively Refining Posterior Approximations

Andrew C. Miller
Harvard University
acm@seas.harvard.edu

Nicholas Foti
University of Washington
nfoti@uw.edu

Ryan P. Adams
Harvard University
rpa@seas.harvard.edu

November 22, 2016

Abstract

We propose a black-box variational inference method to approximate intractable distributions with an increasingly rich approximating class. Our method, termed *variational boosting*, iteratively refines an existing variational approximation by solving a sequence of optimization problems, allowing the practitioner to trade computation time for accuracy. We show how to expand the variational approximating class by incorporating additional covariance structure and by introducing new components to form a mixture. We apply variational boosting to synthetic and real statistical models, and show that resulting posterior inferences compare favorably to existing posterior approximation algorithms in both accuracy and efficiency.

1 Introduction

Variational inference (VI) [2, 16, 30] is a family of methods designed to approximate an intractable *target* distribution (typically known only up to a constant) with a *tractable surrogate* distribution. VI procedures typically minimize the Kullback-Leibler (KL) divergence of the approximation to the target by maximizing an appropriately defined tractable objective. Often, the class of approximating distributions is fixed, and typically excludes the neighborhood surrounding the target distribution, which prevents the variational approximation from becoming arbitrarily close to the true posterior. Often this mismatch between the variational family and the true posterior manifests as underestimating the posterior variances of the model parameters [30].

Markov chain Monte Carlo (MCMC), an alternative class of inference methods, instead approximates target distributions with samples drawn from a Markov chain constructed to leave the target distribution invariant. MCMC methods allow a user to trade computation

time for increased accuracy — drawing more samples will make the approximation closer to the true target distribution. However, MCMC algorithms typically must be run iteratively and it can be difficult to assess convergence to the true posterior. Furthermore, correctly specifying MCMC moves can be more algorithmically restrictive than optimizing an objective (e.g., data subsampling in stochastic gradient methods).

In order to alleviate the mismatch between tractable variational approximations and complicated posterior distributions, we propose a variational inference method that iteratively allows the approximating distribution to become more complex and to eventually represent the true distribution arbitrarily well. This choice allows the practitioner to trade time performing inference against accuracy of posterior estimates, where in the limit the exact posterior can be recovered as in MCMC. Our algorithm grows the complexity of the approximating class in two ways: 1) incorporating richer covariance structure in the component distributions, and 2) by sequentially adding new components to the approximating distribution. Our method builds on black-box variational inference methods using the *re-parameterization trick* [18, 25, 28], applicable to a broad class of target distributions.

The following section discusses variational inference methods, drawing comparisons to alternative approximate inference algorithms. Section 3 and subsections therein describe *variational boosting*. We show how to adapt the re-parameterization trick for mixture approximations in Section 3.1. Section 4 describes various numerical experiments on real and synthetic data.

2 Variational Inference

Given a *target distribution* with density¹ $\pi(x)$ for a random variable $x \in \mathcal{X} \subset \mathbb{R}^d$, variational inference approximates $\pi(x)$ with a *tractable approximate* distribution,² $q(x; \lambda)$, from which we can draw samples and form sample-based estimates of functions of x . Variational methods minimize the KL-divergence, $\text{KL}(q||\pi)$, between $q(\cdot; \lambda)$ and the true π as a function of variational parameters λ [1]. Direct optimization of $\text{KL}(q||\pi)$ is often intractable; however, we can derive a tractable objective based on properties of the KL-divergence. This objective is often referred to as the *evidence lower bound* (ELBO), written

$$\mathcal{L}(\lambda) = \mathbb{E}_{q_\lambda} [\ln \tilde{\pi}(x) - \ln q(x; \lambda)] \quad (1)$$

$$= \mathbb{E}_{q_\lambda} [\ln \pi(x) - \ln q(x; \lambda)] + \ln \mathcal{C} \quad (2)$$

$$= \ln \mathcal{C} - \text{KL}(q_\lambda||\pi) \quad (3)$$

$$\leq \ln \mathcal{C} = \ln \int \tilde{\pi}(x) dx \quad (4)$$

which, due to the positivity of $\text{KL}(q||\pi)$, is a lower bound on the normalization constant³ of $\tilde{\pi}(x)$,

¹We assume $\pi(x)$ is known up to a constant, $\tilde{\pi}(x) = \mathcal{C}\pi(x)$ for some constant \mathcal{C} , omitting \sim to simplify notation.

²We treat the density function as a synecdoche for the entire law, and use $q(x; \lambda)$ and $q_\lambda(x)$ interchangeably at the risk of slight notational abuse.

³Often referred to as the *marginal likelihood*, $p(\text{data})$, in Bayesian inference.

Variational methods typically define (or derive) a family of distributions $Q = \{q(\cdot; \lambda) : \lambda \in \Lambda\}$ parameterized by λ , and maximize the ELBO with respect to $\lambda \in \Lambda$. Most commonly the class Q is fixed, and there exists some (possibly non-unique) $\lambda^* \in \Lambda$ for which $\text{KL}(q||\pi)$ is minimized. When the family Q does not include π , there will be a non-zero KL gap between $q(\cdot; \lambda^*)$ and π , and that discrepancy will realize itself in the form of biased estimates of functions of $x \sim \pi$.

Variational inference is often seen as an alternative to other approximate inference algorithms, most notably Markov chain Monte Carlo (MCMC). MCMC methods construct a Markov Chain such that the target distribution remains invariant (i.e., the target is admitted along the margins). Expectations with respect to the target can be calculated as an average with respect to these correlated samples. MCMC typically enjoys nice asymptotic properties; as the number of samples grows, MCMC samplers represent the true target distribution with higher and higher fidelity. However, rules for constructing *correct* Markov steps are quite restrictive. With a few exceptions [21, 31], most MCMC algorithms require evaluating a log-likelihood that touches all data at each step in the chain (sometimes many times per step). This is problematic in analyses with a large amount of data — MCMC methods are often considered unusable because of this computational bottleneck.

Data sub-sampling, on the other hand, can often be used in conjunction with variational inference methods. Unbiased estimates of the *log-likelihood* based on data sub-sampling can often be used for optimization methods. Because variational methods recast inference as optimization, data sub-sampling can often be a way to make an already efficient approximation even more efficient.

In the next section, we propose an algorithm that iteratively grows the approximating class Q and reframes the VI procedure as a series of optimization problems, resulting in an inference method that can both represent complex distributions and scale to large data sets.

3 Method: Variational Boosting

We define our approximate distribution to be a mixture of C simpler component distributions

$$q^{(C)}(x; \lambda) = \sum_{c=1}^C \rho_c q_c(x; \lambda_c) \quad \text{s.t.} \quad \rho_c \geq 0 \quad \text{and} \quad \sum_c \rho_c = 1 \quad (5)$$

where we have defined component distributions q_c ⁴, mixture component parameters $\lambda = (\lambda_1, \dots, \lambda_C)$, and mixing proportion parameters $\rho = (\rho_1, \dots, \rho_C)$. The component distributions can be any distribution over x from which we can draw samples using a continuous mapping that depends on λ_c (e.g., multivariate normals [14], or a composition of invertible maps [27]).

When posterior expectations and variances are of interest, mixture distributions provide

⁴We denote full mixtures with parenthetical superscripts, $q^{(C)}$, and components with naked subscripts, q_c .

tractable summaries (so long as the component distributions are tractable)⁵. Expectations are easily expressed in terms of component expectations

$$\mathbb{E}_{q^{(C)}}[f(x)] = \int q^{(C)}(x)f(x)dx \quad (6)$$

$$= \sum_c \rho_c \mathbb{E}_{q_c}[f(x)] . \quad (7)$$

In the case of multivariate normal components, the mean and covariance of a mixture are easy to compute in closed form

$$\mathbb{E}_{q^{(C)}}[x] = \sum_c \rho_c \mu(\lambda_c) = \mu^{(C)} \quad (8)$$

$$\mathbb{C}_{q^{(C)}}[x] = \sum_c \rho_c \Sigma(\lambda_c) - \rho_c \left(\mu(\lambda_c) - \mu^{(C)} \right) \left(\mu(\lambda_c) - \mu^{(C)} \right)^\top \quad (9)$$

as are marginal distributions along any set of dimensions

$$q^{(C)}(x_d) = \sum_c \rho_c \mathcal{N}(x_d | \mu_d(\lambda_c), \Sigma_{dd}(\lambda_c)) \quad (10)$$

where $\mu(\lambda_c)$ and $\Sigma(\lambda_c)$ isolate the mean and covariance from variational component parameter λ_c .

Our method begins with a single mixture component, $C = 1$. We use existing BBVI methods to fit the first component parameter, λ_1 , and ρ_1 is fixed to 1 by definition. At the next iteration, we fix λ_1 , and then introduce a new component into the mixture, $q_2(x; \lambda_2)$, a new ELBO objective as a function of new component parameters, λ_2 , and a new mixture weight, ρ_2 . We then optimize this objective with respect to λ_2 and ρ_2 until convergence. At each subsequent iteration, k , we introduce new component parameters and a mixing weight, (λ_k, ρ_k) , which are then optimized according to the new ELBO objective. We refer to this procedure as *variational boosting*, inspired by methods for learning strong classifiers by weighting an ensemble of weak classifiers.

In order for our method to be applicable to a general class of target distributions, we use black-box variational inference methods and the *re-parameterization trick* [18, 25, 28] to fit each component and mixture weights. The re-parameterization trick is a method for obtaining unbiased estimates of the gradient of the ELBO. These gradient estimates can then be used to optimize the ELBO objective using a stochastic gradient optimization method. However, using mixtures as the variational approximation complicates the use of the re-parameterization trick.

⁵Mixtures are also simple to sample from so that more complicated functionals can also easily be estimated.

3.1 The re-parameterization trick and mixture distributions

The re-parameterization trick is a method for computing low-variance estimates of the gradient of an objective for which we only have an unbiased estimator

$$\begin{aligned}\mathcal{L}(\lambda) &= \mathbb{E}_q [\ln \pi(x) - \ln q(x; \lambda)] \\ &\approx \frac{1}{L} \sum_{\ell=1}^L [\ln \pi(x^{(\ell)}) - \ln q(x^{(\ell)}; \lambda)]\end{aligned}$$

where samples $x^{(\ell)}$ are drawn from $q(x; \lambda)$. To obtain a Monte Carlo gradient of $\mathcal{L}(\lambda)$ using the re-parameterization trick, we first separate the randomness needed to generate $x^{(\ell)}$ from the parameters λ , by defining a deterministic map $f_q(x_0; \lambda) = x^{(\ell)}$ such that $x_0 \sim q_0$ implies⁶ $x^{(\ell)} \sim q(x; \lambda)$. Then, we can differentiate through f_q with respect to λ to obtain a gradient estimator.

The re-parameterization trick when q is a mixture, however, is less straightforward. The sampling procedure for a mixture model typically contains a discrete component (i.e., sampling component identities), which is a process that cannot be differentiated through. We circumvent this complication by re-writing the variational objective as a weighted combination of expectations with respect to individual mixture components. Because of the form of the mixture, we can write the ELBO as

$$\begin{aligned}\mathcal{L}(\lambda, \rho) &= \mathbb{E}_q [\ln \pi(x) - \ln q(x; \lambda)] \\ &= \int \left(\sum_{c=1}^C \rho_c q_c(x; \lambda_c) \right) [\ln \pi(x) - \ln q(x; \lambda)] dx \\ &= \sum_{c=1}^C \rho_c \int q_c(x; \lambda_c) [\ln \pi(x) - \ln q(x; \lambda)] dx \\ &= \sum_{c=1}^C \rho_c \mathbb{E}_{q_c} [\ln \pi(x) - \ln q(x; \lambda)]\end{aligned}$$

which is a function of expectations with respect to *mixture components*. If these distributions are continuous, and there exists some function $f_c(x_0; \lambda)$ such that $x = f_c(x_0; \lambda)$ and $x \sim q_c(\cdot; \lambda)$ when $x_0 \sim q_0$, then we can apply the re-parameterization trick to each component to obtain gradients of the ELBO

$$\begin{aligned}\nabla_{\lambda_c} \mathcal{L}(\lambda, \rho) &= \nabla_{\lambda_c} \sum_{c=1}^C \rho_c \mathbb{E}_{x \sim q(x; \lambda)} [\ln \pi(x) - \ln q(x; \lambda)] \\ &= \sum_{c=1}^C \rho_c \mathbb{E}_{x_0 \sim q_0} [\nabla_{\lambda_c} \ln \pi(f_c(x_0; \lambda_c)) - \nabla_{\lambda_c} \ln q(f_c(x_0; \lambda_c))].\end{aligned}$$

Variational Boosting uses the above fact with the re-parameterization trick in a component-by-component manner, allowing us to improve the variational approximation as we incorporate and fit new components.

⁶Here, q_0 is some base distribution that is, importantly, *not* a function of λ .

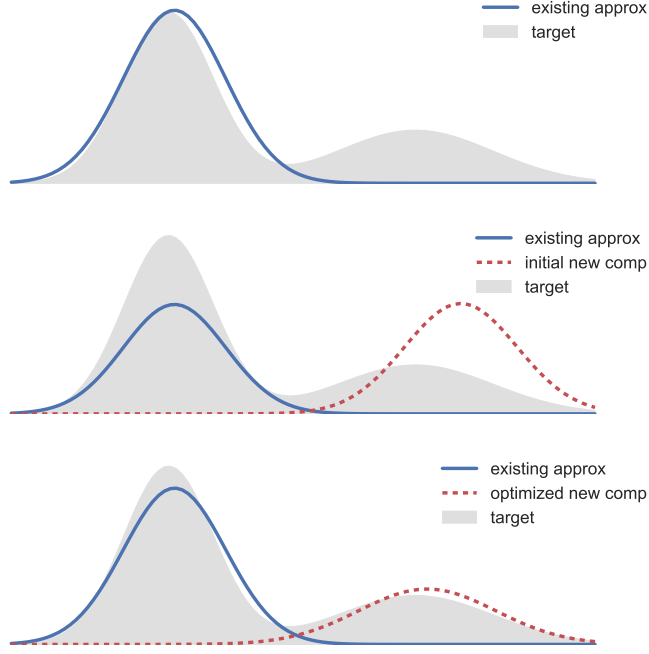


Figure 1: Illustrative one-dimensional example of adding a new component in *variational boosting*. Top: initial approximation with a single component (solid green). Middle: a new component (dotted red), is initialized using Algorithm 1. Bottom: the new component parameters and mixing weights are optimized using Monte Carlo gradients of the ELBO. Note that this allows the mass of the existing components to rise and fall, but not shift in space.

3.2 Adding Components

In this section we present details of the proposed algorithm. We first describe the process of fitting a single component and then the process for adding an additional component to an existing mixture distribution.

Fitting the first component The procedure starts by fitting an approximation to $\pi(x)$ with a distribution that consists of a single component. We do this by maximizing the first ELBO objective

$$\mathcal{L}^{(1)}(\lambda_1) = \mathbb{E}_q [\ln \pi(x) - \ln q_1(x; \lambda_1)] \quad (11)$$

$$\lambda_1^* = \arg \max_{\lambda_1} \mathcal{L}^{(1)}(\lambda_1). \quad (12)$$

Depending on the forms of π and q_1 , optimizing the ELBO can be accomplished by various methods. One general method for fitting a continuous valued component is to compute stochastic, unbiased gradients of $\mathcal{L}(\lambda_1)$, and use stochastic gradient optimization. After convergence (or close to it) we fix λ_1 to be λ_1^* .




Figure 2: Sequence of increasing complex approximate posteriors, with $C = 1, 2, 3, 4$. The background (grey/black) contours depict the target distribution, and the foreground (red) contours depict the iterative approximations.

Fitting component $C + 1$ After iteration C , our current approximation to $\pi(x)$ is a mixture distribution with C components

$$q^{(C)}(x; \lambda) = \sum_{c=1}^C \rho_c q_c(x; \lambda_c) \quad (13)$$

where $\lambda = (\{\rho_c, \lambda_c\}_c)$ is a list of component parameters and mixing weights, and $q_c(x; \lambda_c)$ is the component distribution parameterized by λ_c . Adding a new component introduces a new component parameter, λ_{C+1} , and a new mixing weight, ρ_{C+1} . In this section, the mixing parameter $\rho_{C+1} \in [0, 1]$ mixes between the new component, $q_{C+1}(\cdot; \lambda_{C+1})$ and the existing approximation, $q^{(C)}$. The new approximate distribution is

$$q^{(C+1)}(x; \rho_{C+1}, \lambda_{C+1}) = (1 - \rho_{C+1})q^{(C)}(x) + \rho_{C+1}q_{C+1}(x; \lambda_{C+1}). \quad (14)$$

The new optimization objective, as a function of ρ_{C+1} and λ_{C+1} is

$$\mathcal{L}^{(C+1)}(\rho_{C+1}, \lambda_{C+1}) = \mathbb{E}_{x \sim q^{(C+1)}} \left[\ln \pi(x) - \ln q^{(C+1)}(x; \lambda_{C+1}, \rho_{C+1}) \right] \quad (15)$$

$$= (1 - \rho_{C+1}) \mathbb{E}_{q^{(C)}} \left[\ln \pi(x) - \ln q^{(C+1)}(x; \lambda_{C+1}, \rho_{C+1}) \right] \quad (16)$$

$$+ \rho_{C+1} \mathbb{E}_{q_{C+1}} \left[\ln \pi(x) - \ln q^{(C+1)}(x; \lambda_{C+1}, \rho_{C+1}) \right]. \quad (17)$$

Above we have separated out two expectations — one with respect to the existing approximation (which is fixed), and the other with respect to the new component distribution. Because we have fixed the existing approximation, we only need to optimize the new component parameters, $\lambda_{C+1}, \rho_{C+1}$, allowing us to use the re-parameterization trick to obtain gradients of $\mathcal{L}^{(C+1)}$. As we have fixed the existing component distribution and we only need to optimize the new component λ_{C+1} , we can use the re-parameterization trick and Monte Carlo gradients to optimize $\mathcal{L}^{(C+1)}$ with respect to ρ_{C+1} and λ_{C+1} .

Figure 1 illustrates the algorithm on a simple one-dimensional example — showing the initialization of a new component and the resulting mixture after optimizing the second objective, $\mathcal{L}^{(2)}(\rho_2, \lambda_2)$. Figure 2 depicts the result of the variational boosting procedure on a two-dimensional, multi-modal target distribution. In both cases, the component distributions are Gaussians with diagonal covariance.

3.3 Structured Multivariate Normal Components

Though our method can use any component distribution that can be sampled using a continuous mapping, a sensible choice of component distribution is a multivariate normal

$$q(x; \lambda) = \mathcal{N}(x; \mu(\lambda), \Sigma(\lambda)) \quad (18)$$

$$= |2\pi\Sigma(\lambda)|^{-1/2} \exp\left(-\frac{1}{2}(x - \mu(\lambda))^\top \Sigma(\lambda)^{-1}(x - \mu(\lambda))\right) \quad (19)$$

where the variational parameter λ is transformed into a mean vector $\mu(\lambda)$ and covariance matrix $\Sigma(\lambda)$.

Specifying the structure of the covariance matrix is a choice that largely depends on the dimensionality of x ($x \in \mathbb{R}^D$) and correlation structure of the target distribution. A common first-choice of covariance parameterization is a diagonal matrix

$$\Sigma(\lambda) = \text{diag}(\sigma_1^2, \dots, \sigma_D^2) \quad (20)$$

which implies that x is independent across dimensions. When the approximation only consists of one component, this structure is commonly referred to as the *mean field* family. While computationally efficient, mean field approximations cannot model posterior correlations, which often leads to underestimation of marginal variances. Additionally, when diagonal covariances are used as the component distributions in Eq. (5) the resulting mixture may require a exorbitant number of components to represent the strong correlations. Further, the independence restriction can introduce local optima in the variational objective [30].

On the other end of the spectrum, we can parameterize the entire covariance matrix by parameterizing the lower triangle of a Cholesky decomposition, L , such that $LL^\top = \Sigma$. This allows Σ to be any positive semi-definite matrix, enabling q to have the full flexibility of a D -dimensional multivariate normal distribution. However, this introduces $D(D+1)/2$ parameters, which can become computationally cumbersome when D is large. Furthermore, it may not be the case that all pairs of variables exhibit posterior correlation, particularly in multi-level models where different parameter types may be more or less independent in the posterior.

Alternatively, we can incorporate *some* capacity to capture correlations between dimensions of x without introducing many more parameters. The next subsection discusses

a covariance specification that provides this tradeoff, while remaining computationally tractable within the BBVI framework.

Low-rank plus diagonal covariance Black-box variational inference methods with the re-parameterization trick rely on sampling from the variational distribution, and efficiently computing (or approximating) the entropy of the variational distribution. For multivariate normal distributions, the entropy is a function of the determinant of the covariance matrix, Σ , while computing the log likelihood requires inverting the covariance matrix, Σ^{-1} . When the dimensionality of the target, D , is large, computing determinants and inverses will be $O(D^3)$ and therefore may be prohibitively expensive to compute at every iteration.

However, it may be unnecessary to represent all $D(D-1)/2$ possible correlations in the target distribution, particularly if certain dimensions are close to independent. One way to increase the capacity of $q(x; \lambda)$ is to model the covariance as a *low-rank plus diagonal* (LR+D) matrix

$$\Sigma = CC^\top + \text{diag}(\exp(v)) \quad (21)$$

where $C \in \mathbb{R}^{D \times r}$ and $v \in \mathbb{R}^D$ are the rank- r and log diagonal components. Note that both C and v are represented by the component parameter λ .

The choice of r presents a tradeoff — with a larger rank, the variational approximation can be more flexible; with a lower rank, the computations necessary for fitting the variational approximation can be more efficient. As a concrete example, in the Experiments section we present a $D = 40$ dimensional posterior resulting from a non-conjugate hierarchical model, and we show that a “rank $r = 2$ plus diagonal” covariance does an excellent job capturing all $D(D-1)/2 = 780$ pairwise correlations and D marginal variances. Incorporating more components using the variational boosting framework further improves the approximation of the distribution.

To use the re-parameterization trick with this low rank covariance, we can simulate from q in two steps

$$z^{(lo)} \sim \mathcal{N}(0, I_r) \quad (22)$$

$$z^{(hi)} \sim \mathcal{N}(0, I_D) \quad (23)$$

$$x = Cz^{(lo)} + \mu + \mathcal{I}(v/2)z^{(hi)} \quad (24)$$

where $z^{(lo)}$ generates the randomness due to the low-rank structure, and $z^{(hi)}$ generates the randomness due to the diagonal structure. We use the operator $\mathcal{I}(a) = \text{diag}(\exp(a))$ for notational brevity. This generative process can be differentiated through, yielding Monte Carlo estimates of the gradient with respect to C and v suitable for stochastic optimization.

In order to use LR+D covariance structure within variational boosting, we will need to efficiently compute the determinant and inverse of Σ . The matrix determinant lemma [11]

allows us to represent the determinant of Σ as the product of two determinants

$$|CC^\top + \mathcal{I}(v)| = |\mathcal{I}(v)| |I_r + C^\top \mathcal{I}(-v)C| \quad (25)$$

$$= \exp\left(\sum_d v_d\right) |I_r + C^\top \mathcal{I}(-v)C| \quad (26)$$

where the left term is simply the product of the diagonal component, and the right term is the determinant of a $r \times r$ dimensional matrix, computable in $O(r^3)$ time.

Similarly, the Woodbury matrix identity [9] allows us to represent the inverse of Σ as

$$(CC^\top + \mathcal{I}(v))^{-1} = \mathcal{I}(-v) - \mathcal{I}(-v)C(I_r + C^\top \mathcal{I}(-v)C)^{-1}C^\top \mathcal{I}(-v) \quad (27)$$

which involves the inversion of a smaller, $r \times r$ matrix, which can be done in $O(r^3)$ time. Importantly, the above operations are efficiently differentiable and amenable for use in the BBVI framework.

Fitting the rank To specify the ELBO objective, we need to choose a rank r for the component covariance. Because fitting a single component is relatively cheap, we start by a single component with rank $r = 0$, continue to fit $r = 1, 2, \dots$, and rely on a heuristic stopping criterion. For a single Gaussian, one such criterion is the average change in marginal variances — if the marginal variation along each dimension remains the same from rank r to $r + 1$, then the new covariance component is not incorporating explanatory power, particularly if marginal variances are of interest. As the $KL(q||\pi)$ objective tends to underestimate variances when restricted to a particular model class, we observe that the marginal variances grow as new covariance rank components are added. When fitting rank $r + 1$, we can monitor the average absolute change in marginal variance (or standard deviation) as more covariance structure is incorporated. Figure 9 in Section 4 depicts this measurement for a $D = 37$ -dimensional posterior.

3.4 Initializing Components

Introducing a new component requires initialization of component parameters. When our component distributions are mixtures of Gaussians, we found that the optimization procedure is sensitive to initialization. This section describes an importance-weighting scheme for initialization that produces (empirically) good initial values of component and mixing parameters.

Conceptually, a good initial component is located in a region of the target $\pi(x)$ that is underrepresented by the existing approximation $q^{(C)}$. A good initial weight is close to the proportion of mass in the unexplained region. Following this principle, we construct this component by first drawing importance-weighted samples from our existing approximation

$$x^{(\ell)} \sim q^{(C)}, \quad w^{(\ell)} = \frac{\pi(x^{(\ell)})}{q^{(C)}(x^{(\ell)})} \quad \text{for } \ell = 1, \dots, L. \quad (28)$$

The samples with the largest weights $w^{(\ell)}$ tell us where regions of the target are unexplained by our approximation. In fact, as L grows, and if $q^{(C)}$ is “close” enough

to π , we can interpret $\{x^{(\ell)}, w^{(\ell)}\}$ as a weighted sample from π . Based on this interpretation, we can fit a mixture distribution (or *some components* of a mixture distribution) to this weighted sample using maximum likelihood, and recover a type of target approximation. For mixture distributions, an efficient inference procedure is Expectation-Maximization (EM) [3].

This approach, however, presents a few complications. First, we must adapt EM to fit a *weighted* sample. Second, importance weights can suffer from extremely high variance — one or two $w^{(\ell)}$ values may be extremely large compared to all other weights. This destabilizes our new component parameters and mixing weight, particularly the variance of the component. Intuitively, if a single weight $w^{(\ell)}$ is extremely large, this would correspond to many samples being located in a single location, and maximum likelihood with EM would want to shrink the variance of the new component to zero right on that location. To combat this behavior, we use a simple method to break up the big weights using a resampling and re-weighting step before applying weighted EM. Empirically, this improves our new component initializations and subsequent ELBO convergence.

Weighted EM Expectation-maximization is typically used to perform maximum likelihood in latent variable models. Mixture distributions are easily represented with latent variables — a sample’s latent variable corresponds to the mixture component that produced it. EM starts with some initialization of model parameters (e.g., component means, variances and mixing weights). The algorithm then iterates between two steps: 1) the *E-step*, which computes the distribution over the latent variables given the current setting of parameters, and 2) the *M-step*, which maximizes the *expected complete data log-likelihood* with respect to the distributions computed in the E-step.

We suppress details of the general treatment of EM, and focus on EM for mixture models as presented in [1]. For mixture distributions, the E-step computes “responsibilities”, or the probability that a datapoint came from one of the components. The M-step then computes a weighted maximum likelihood, where the log-likelihood of a datapoint for a particular component is weighted by the associated “responsibility”. This weighted maximum likelihood is an easy entry-point for an additional set of weights — weights associated with each datapoint from the importance-weighting.

More concretely, for a sample of data, $x^{(\ell)}$, C mixture components, and current mixture component parameters and weights $\lambda = \{\rho_c, \lambda_c\}_{c=1}^C$, the E-step computes the following quantities

$$\gamma_c^{(\ell)} = p(z^{(\ell)} = c | x^{(\ell)}, \lambda) \quad (29)$$

$$\propto p(x^{(\ell)} | z^{(\ell)}, \lambda_c = c) p(z^{(\ell)} = c) \quad (30)$$

where $\gamma_c^{(\ell)}$ is the “responsibility” of cluster c for datapoint ℓ . The M-step then computes component parameters by a weighted maximum likelihood

$$\lambda_c^* = \arg \max_{\lambda} \sum_{\ell=1}^L \gamma_c^{(\ell)} \cdot \ln p(x^{(\ell)} | z^{(\ell)} = c, \lambda_c). \quad (31)$$

Algorithm 1 Importance-weighted initialization of new components. This algorithm takes in the target distribution, $\pi(x)$, the current approximate distribution $q^{(C)}(x)$, and a number of samples L . This returns an initial value of new component parameters, λ_{C+1} and a new mixing weight ρ_{C+1} .

```

1: procedure INITCOMP( $\pi, q^{(C)}, L$ )
2:    $x^{(\ell)} \sim q^{(C)}$  for  $\ell = 1, \dots, L$  ▷ sample from existing approx
3:    $w^{(\ell)} \leftarrow \frac{\pi(x^{(\ell)})}{q^{(C)}(x^{(\ell)})}$  ▷ set importance weights
4:    $\mathcal{O} \leftarrow \text{outlier-weights}(\{w^{(\ell)}\})$ 
5:    $q^{(IW)} \leftarrow \text{make-mixture}(\mathcal{O}, \{w^{(\ell)}\}, q^{(C)})$  ▷ break up big weights
6:    $x_r^{(\ell)} \sim q^{(IW)}$  for  $\ell = 1, \dots, L$  ▷ sample from new mixture
7:    $w_r^{(\ell)} \leftarrow \frac{\pi(x_r^{(\ell)})}{q^{(IW)}(x_r^{(\ell)})}$  ▷ re-sampled importance weights
8:    $\lambda_{C+1}, \rho_{C+1} \leftarrow \text{weighted-em}(\{x_r^{(\ell)}, w_r^{(\ell)}\})$  ▷ fit new component
9:   return  $\lambda_{C+1}, \rho_{C+1}$ 

```

To incorporate importance weights $w^{(\ell)}$, we only need to slightly change the M-step.

$$\lambda_c^* = \arg \max_{\lambda} \sum_{\ell=1}^L w^{(\ell)} \cdot \gamma_c^{(\ell)} \cdot \ln p(x^{(\ell)} | z^{(\ell)} = c, \lambda_c) \quad (32)$$

Because we are adding a new component, we would like our weighted EM routine to leave the remaining components unchanged. For instance, we want $\lambda_1, \dots, \lambda_{C-1}$ to be fixed, while λ_C is free to explain the weighted sample. This can be accomplished in a straightforward manner by simply clamping the first $C - 1$ parameters during the M-step.

Resampling importance weights If our current approximation $q^{(C)}$ is sufficiently different in certain regions of the posterior, then some weights $w^{(\ell)}$ will end up being large compared to other weights. For instance, the objective $\text{KL}(q||p)$ tends to under-cover regions of the posterior, allowing $\pi(x)$ to be much larger than $q^{(C)}(x)$, meaning the weight associated with x will be large. This will create instability in the weighted EM approximation — likelihood maximization will want to put a zero-variance component on the single highest-weighted sample, which does not accurately reflect the local curvature of $\pi(x)$. To combat this, we construct a slightly more complicated proposal distribution. Conceptually, we first create this naïve importance-weighted sample, and then find samples with outlier weights, and break those samples up. We do this by constructing a new proposal distribution that mixes the existing proposal, $q^{(C)}$, and component means located at the outlier samples. We define this proposal to be

$$q^{(IW)}(x) = p_0 q^{(C)}(x) + \sum_{\ell \in \mathcal{O}} w^{(\ell)} \mathcal{N}(x | x^{(\ell)}, \Sigma^{(\ell)}) \quad (33)$$

where $\ell \in \mathcal{O}$ denote the set of outlier samples from our original sample, and $p_0 = 1 - \sum_{\ell \in \mathcal{O}} w^{(\ell)}$ is the mass not placed on outlier samples. The variance of each outlier component, $\Sigma^{(\ell)}$ is set to some heuristic value — we typically use the diagonal of the covariance of $q^{(C)}$ as a good-enough guess.

We then create a new importance-weighted sample, using $q^{(IW)}$ and $\pi(x)$ just as we did before. By placing new components (with some non-zero variance) on the outlier samples, which are known to be in a region of high target probability and low approximate probability, we assume that there is more local probability around that region that needs to be explored. This allows us to inflate the local variance of the samples in this region — the region that weighted EM will place a component. Algorithm 1 unites the components from above sections into our final initialization procedure.

3.5 Related Work

Using a mixture model as an approximating distribution in variational inference is a well-studied idea. Mixtures of mean field approximations [14] introduced mean field-like updates for a mixture approximation using a bound on the entropy term and model-specific parameter updates. Nonparametric variational inference [8] is a black-box variational inference algorithm that approximates a target distribution with a mixture of equally-weighted isotropic normals. The authors use a lower-bound on the entropy term in the ELBO to make the optimization procedure tractable. Similarly, [28] present a method for fitting mixture distributions as an approximation. However, their method is restricted to mixture component distributions within the exponential family, and a joint optimization procedure. Finally, we note that [10] independently and in parallel proposed a closely-related idea for iterative “boosted” construction of variational mixture approximations.

Using a low-rank Gaussian as a variational approximation was explored in [29], using a PCA-like algorithm. We fit the low-rank components of a Gaussian using black-box methods and joint optimization.

We also note that mixture distributions are a type of hierarchical variational model [26], where the component identity can be thought of as latent variables in our variational distribution. While in [26] the authors optimize a lower bound on the ELBO to fit general hierarchical variational models, our approach integrates out the discrete latent variables because it is tractable to do so.

4 Experiments and Analysis

To supplement the illustrative synthetic examples, in this section we apply variational boosting to approximate various intractable posterior distributions resulting from real statistical analyses.

4.1 Hierarchical Binomial Regression

We test out our posterior approximation on a hierarchical binomial regression model.⁷ We borrow an example from [4], and estimate the binomial rates of success (batting

⁷Model and data from the `mc-stan` case studies, <http://mc-stan.org/documentation/case-studies/pool-binary-trials.html>




Figure 3: Comparison of univariate and bivariate marginals for the binomial hierarchical model. Each histogram/scatterplot results from 20,000 NUTS samples. Top left: Bivariate marginal (κ, θ_0) HMC samples and a mean field approximation (MFVI). Top Right, the same bivariate marginal, and the Variational Boosting approximation. Bottom: comparison of NUTS, MFVI, and VBoost on univariate marginals (global parameters).

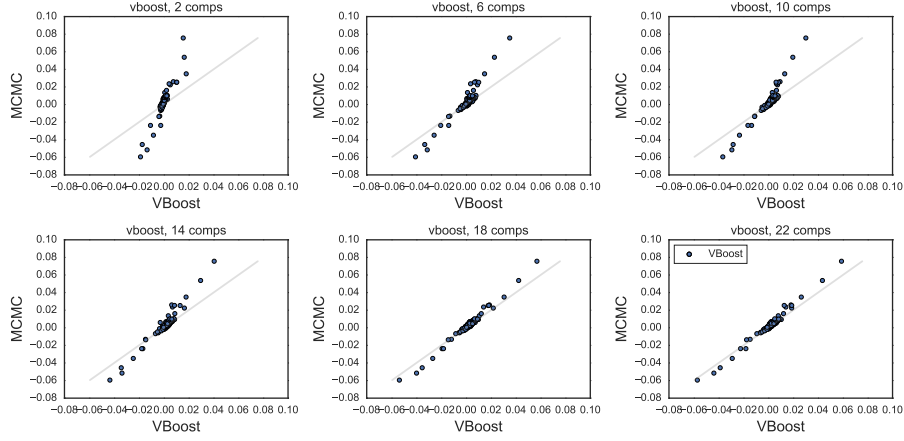


Figure 4: Comparison of posterior covariances for the $D = 20$ -dimensional **baseball** model (hierarchical binomial regression). Each plot corresponds to the variational boosting algorithm as it incorporates new diagonal-covariance mixture components. The vertical axis indicates covariance estimates from 20k samples of NUTS (shared across all plots). We see that as more components are added, the variational boosting covariance estimates more closely match the MCMC covariance estimates.

averages) of baseball players using a hierarchical model. The model describes a latent

“skill” parameter for baseball players — the probability of obtaining a hit in a given at bat. The model of the data is

$\phi \sim \text{Unif}$	hyper prior
$\kappa \sim \text{Pareto}(1, 1.5)$	hyper prior
$\theta_j \sim \text{Beta}(\phi \cdot \kappa, (1 - \phi) \cdot \kappa)$	player j prior
$y_j \sim \text{Binomial}(K_j, \theta_j)$	player j hits

where y_j is the number of successes (hits) player j has attempted in K_j attempts (at bats). Each player has a latent success rate θ_j , which is governed by two global variables κ and ϕ . There are 18 players in this example, creating a posterior distribution with $D = 20$ parameters. This model is not conjugate, and requires approximate Bayesian inference.

We use `adam` [17] for each stochastic optimization problem with default parameters. For stochastic gradients, we use 400 samples for the new component, and 400 samples for the previous component. In all experiments, we use `autograd` [23, 22] to obtain automatic gradients with respect to new component parameters.

To highlight the fidelity of our method, we compare Variational Boosting to mean field VI and the No-U-Turn Sampler (NUTS) [13]. The empirical distribution resulting from 20k NUTS samples is considered the “ground truth” posterior in this example. Figure 3 compares a selection of univariate and bivariate posterior marginals. We see that Variational Boosting is able to closely match the NUTS posteriors, improving upon the MFVI approximation.

Figure 4 compares the variational boosting covariance estimates to the “ground truth” estimates of MCMC at various stages of the algorithm.

4.2 Multi-level Poisson GLM

We apply variational boosting to approximate the posterior for a common hierarchical model, a hierarchical Poisson GLM. This model was formulated to measure the relative rates of stop-and-frisk events for different ethnicities and in different precincts [6], and has been used as illustrative example of multi-level modeling [7, Chapter 15, Section 1].

The model incorporates a precinct and ethnicity effect to describe the relative rate of stop-and-frisk events.

$\mu \sim \mathcal{N}(0, 10^2)$	mean offset
$\ln \sigma_\alpha^2, \ln \sigma_\beta^2 \sim \mathcal{N}(0, 10^2)$	group variances
$\alpha_e \sim \mathcal{N}(0, \sigma_\alpha^2)$	ethnicity effect
$\beta_p \sim \mathcal{N}(0, \sigma_\beta^2)$	precinct effect
$\ln \lambda_{ep} = \mu + \alpha_e + \beta_p + \ln N_{ep}$	log rate
$Y_{ep} \sim \mathcal{P}(\lambda_{ep})$	stop-and-frisk events

where Y_{ep} are the number of stop-and-frisk events within ethnicity group e and precinct p over some fixed period of time; N_{ep} is the total number of arrests of ethnicity group e in precinct p over the same period of time; α_e and β_p are the ethnicity and precinct effects.




Figure 5: Comparison of single Gaussian component marginals by rank for a 37-dimensional Poisson GLM posterior. The top left plot is a diagonal Gaussian approximation. The next plots show how the marginal variances inflate as the covariance is allotted more capacity.

As before, we simulate 20k NUTS samples, and compare various variational approximations. Because of the high posterior correlations present in this example, variational boosting with diagonal covariance components is inefficient in its representation of this structure. As such, this example more heavily relies on the low-rank approximation to shape the posterior.

Figure 5 shows how increasing the rank of a single multivariate normal component can result in better variance approximations. Figure 6 shows a handful of bivariate marginal posterior approximations as a function of covariance rank. Figure 7 shows the same bivariate marginals as more rank-3 components are added to the approximation. Lastly, Figure 8 compares the marginal standard deviations and covariances to MCMC-based measurements. These results indicate that while the incorporation of covariance structure increases the accuracy of marginal variance approximations, the non-Gaussianity afforded by the incorporation of mixture components allows for a better posterior approximation that translates into even more accurate moment estimates.

5 Discussion and Conclusion

We have proposed a practical variational inference method that incorporates new components into the approximation and is applicable to a large number of Bayesian models of interest. We demonstrated the ability of the method to learn rich representations of complex posteriors over a moderate number of parameters.

We see a few avenues of future work. First, while it is known that mixtures of Gaus-




Figure 6: A sampling of bivariate marginals for a single Gaussian component marginals by rank for a $D = 37$ -dimensional Poisson GLM posterior. Although there are a total of 666 covariances to be approximated, only a few directions in the D -dimensional parameter space exhibit non-trivial correlations.

sians can approximate smooth distributions to arbitrary precision (with enough components) [5], it remains an open question if our approach of fixing and iteratively adding components using this sequence of ELBO objectives will converge. Existing work has shown that this is the case for the alternative direction of KL-divergence, $\text{KL}(\pi||q)$ [19, 24], but it remains to be shown for $\text{KL}(q||\pi)$. Furthermore, the rate of




Figure 7: A sampling of bivariate marginals for mixtures of rank-3 Gaussians at various stages of the variational boosting algorithm for the $D = 37$ -dimensional **frisk** model. Introducing new mixture components allows the posterior to take a non-Gaussian shape, most exhibited in the third column.

convergence would, ideally, be characterized.

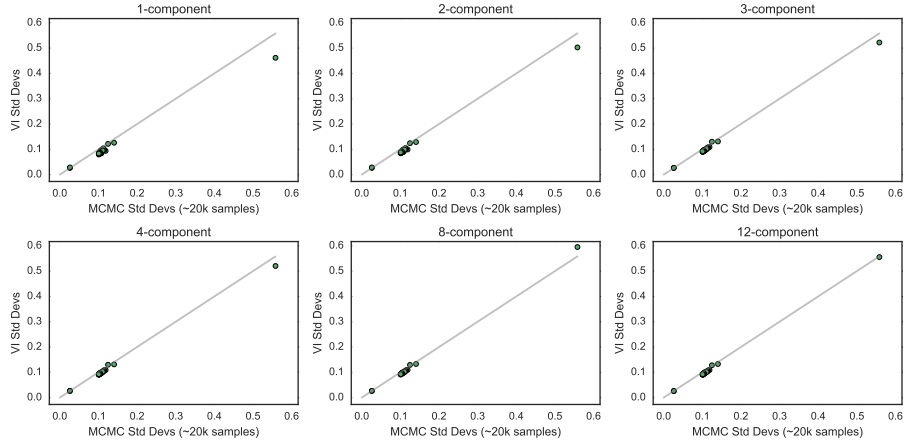
The variational boosting framework allows for more flexible component distributions. For instance, compositions of invertible maps have been used to enrich variational families [27], as well as auxiliary variable variational models [20].

Although our mixture component fitting algorithm is greedy, our determination of an appropriate covariance rank is not. We imagine that we can use the result of the $r = 0$ diagonal covariance to inform the procedure for $r = 1$, and so on. We leave this sort of *nested boosting* for low-rank determination to future work.

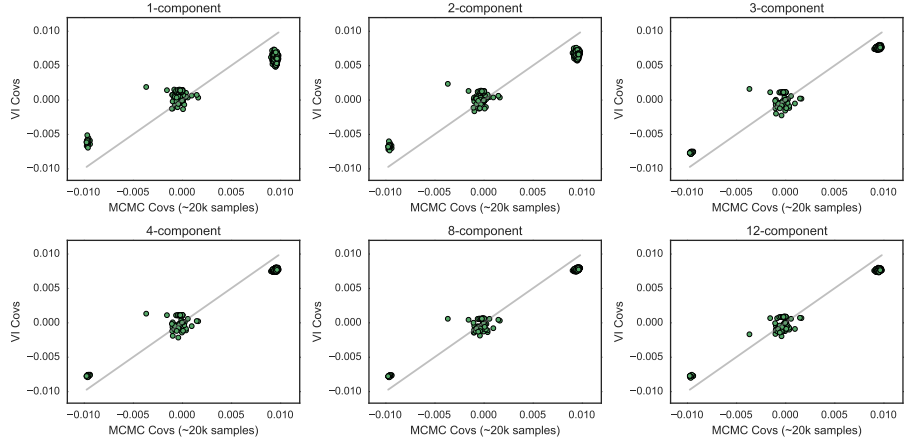
When optimizing parameters of a variational family, it has been shown that the natural gradient can be more robust and lead to better optima [12, 15]. While the Fisher information, required for computing the natural gradient, for a single Gaussian component can be computed in closed form, it is less straightforward for a mixture component. We hope to incorporate natural gradient updates in future work.

Acknowledgments

The authors would like to acknowledge Arjumand Masood, Mike Hughes, and Finale Doshi-Velez for helpful conversations. ACM is supported by the Applied Mathematics Program within the Office of Science Advanced Scientific Computing Research of the U.S. Department of Energy under contract No. DE-AC02-05CH11231. NF is supported by a Washington Research Foundation Innovation Postdoctoral Fellowship in Neuroengineering



(a) Marginal standard deviations



(b) Pairwise covariances

Figure 8: A comparison of standard deviations and covariances for the **frisk** model. The MCMC-inferred values are along the horizontal axis, with the variational boosting values along the vertical axis. While the rank 3 plus diagonal covariance structure is able to account for most of the marginal variances, the largest one is still underestimated. Incorporating more rank 3 components allows the approximation to account for this variance. Similarly, the non-zero covariance measurements improve as more components are added.

and Data Science. RPA is supported by NSF IIS-1421780 and the Alfred P. Sloan Foundation.

References

- [1] C Bishop. Pattern recognition and machine learning, 2006.




Figure 9: Mean percent change in marginal variances for the Poisson GLM. After rank 5, the average percent change is less than 5% — this estimate is slightly noisy due to the stochastic optimization procedure.

- [2] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *arXiv preprint arXiv:1601.00670*, 2016.
- [3] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- [4] Bradley Efron and Carl Morris. Data analysis using stein’s estimator and its generalizations. *Journal of the American Statistical Association*, 70(350):311–319, 1975.
- [5] V. A. Epanechnikov. Non-parametric estimation of a multivariate probability density. *Theory Probab. Appl.*, 14(1):153–158, 1967.
- [6] Andrew Gelman, Jeffrey Fagan, Alex Kiss, et al. An analysis of the nypd’s stop-and-frisk policy in the context of claims of racial bias. *Journal of the American Statistical Association*, 102:813–823, 2007.
- [7] Andrew Gelman and Jennifer Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, 2006.
- [8] Samuel Gershman, Matt Hoffman, and David M Blei. Nonparametric variational inference. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 663–670, 2012.
- [9] G. H. Golub and C. F. Van Loan. *Matrix Computations*. JHU Press, 2013.
- [10] Fangjian Guo, Xiangyu Wang, Kai Fan, Tamara Broderick, and David B. Dunson. Boosting variational inference. *arXiv:1611.05559 [stat.ML]*, 2016.
- [11] D. A. Harville. *Matrix Algebra from a Statistician’s Perspective*. Springer-Verlag, 1997.
- [12] Matthew D Hoffman, David M Blei, Chong Wang, and John William Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- [13] Matthew D Hoffman and Andrew Gelman. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [14] Tommi S Jaakkola and Michael I Jordan. Improving the mean field approximation via the use of mixture distributions. In *Learning in graphical models*, pages 163–173. Springer, 1998.

- [15] Matthew J. Johnson, David K. Duvenaud, Alex B. Wiltschko, Sandeep R. Datta, and Ryan P. Adams. Composing graphical models with neural networks for structured representations and fast inference. *Arxiv preprint arXiv:1603.06277*, 2016.
- [16] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- [17] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [19] Q. J. Li and A. R. Barron. Mixture density estimation. In *Advances in Neural Information Processing Systems*, 1999.
- [20] Lars Maaløe, Casper Kaae Sønderby, Søren Kaae Sønderby, and Ole Winther. Auxiliary deep generative models. *arXiv preprint arXiv:1602.05473*, 2016.
- [21] Dougal Maclaurin and Ryan P Adams. Firefly monte carlo: Exact mcmc with subsets of data. *arXiv preprint arXiv:1403.5693*, 2014.
- [22] Dougal Maclaurin, David Duvenaud, and Ryan P. Adams. Autograd: Reverse-mode differentiation of native python. *ICML workshop on Automatic Machine Learning*, 2015.
- [23] Dougal Maclaurin, David Duvenaud, Matthew Johnson, and Ryan P. Adams. Autograd: Reverse-mode differentiation of native Python, 2015.
- [24] A. Rakhlin, Panchenko D., and Mukherjee S. Risk bounds for mixture density estimation. *ESAIM: Probability and Statistics*, 9:220–229, 2006.
- [25] Rajesh Ranganath, Sean Gerrish, and David M Blei. Black box variational inference. In *AISTATS*, pages 814–822, 2014.
- [26] Rajesh Ranganath, Dustin Tran, and David M Blei. Hierarchical variational models. In *International Conference on Machine Learning*, 2016.
- [27] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1530–1538, 2015.
- [28] Tim Salimans, David A Knowles, et al. Fixed-form variational posterior approximation through stochastic linear regression. *Bayesian Analysis*, 8(4):837–882, 2013.
- [29] M. W. Seeger. Gaussian covariance and scalable variational inference. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.
- [30] Martin J Wainwright and Michael I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.
- [31] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.